arXiv:1701.07166v1 [cs.DC] 25 Jan 2017

Personalized Classifier Ensemble Pruning Framework for Mobile Crowdsourcing

Shaowei Wang[1], Liusheng Huang[2], Pengzhan Wang[3],
Hongli Xu[4], and Wei Yang[5]

School of Computer Science and Technology,
University of Science and Technology of China, Hefei 230027, Anhui, China
{wangsw[1], pzwang[3]}@mail.ustc.edu.cn,
{lshuang[2], xuhongli[4], qubit[5]}@ustc.edu.cn

Abstract. Ensemble learning has been widely employed by mobile applications, ranging from environmental sensing to activity recognitions. One of the fundamental issue in ensemble learning is the trade-off between classification accuracy and computational costs, which is the goal of ensemble pruning. During crowdsourcing, the centralized aggregator releases ensemble learning models to a large number of mobile participants for task evaluation or as the crowdsourcing learning results, while different participants may seek for different levels of the accuracy-cost trade-off. However, most of existing ensemble pruning approaches consider only one identical level of such trade-off. In this study, we present an efficient ensemble pruning framework for personalized accuracy-cost trade-offs via multi-objective optimization. Specifically, for the commonly used linear-combination style of the trade-off, we provide an objective-mixture optimization to further reduce the number of ensemble candidates. Experimental results show that our framework is highly efficient for personalized ensemble pruning, and achieves much better pruning performance with objective-mixture optimization when compared to state-of-art approaches.

Keywords: Ensemble Pruning, Mobile Crowdsourcing, Bi-objective Optimization, Pareto Optimization

1 Introduction

A classifier ensemble leverages a bundle of diverse base classifiers (e.g., SVM, k-NN, Bagging classifier, Boosting classifier) to complete classification tasks by aggregating the classification results of the base classifiers. Since individual classifiers may perform poorly whereas the ensemble of these classifiers usually significantly improves, classifier ensembles are popular and are widely adopted in mobile device applications, ranging from environment-centric scenarios (e.g., place classification[18]) to people-centric scenarios (e.g., activity recognition [3], human fall detection [7][14]).

Though classifier ensemble could potentially improve classification accuracy, evaluating every member in the ensemble is much more expensive than the evaluation of an individual classifier. On the other hand, the number of classifiers in the ensemble limits the achievable classification accuracy. Such dilemma calls for the trade-off between classification accuracy and the classification cost (e.g., the size of the ensemble or time and memory consumptions of the ensemble) for classifier ensembles. Hence, *ensemble pruning* [17] has been proposed to select a subset of all possible classifiers meanwhile avoid harming classification accuracy. It's worth noting that the accuracy-cost trade-off is especially important in mobile crowdsourcing scenarios, where mobile devices' power saving is a crucial issue and realtime online classification is usually needed.

Specifically, in the crowdsourcing scenarios, the centralized crowdsourcing server may need to publish pruned classifier ensembles to every crowdsourcing participants, mainly due to two types of purposes. One purpose is facilitating each participant to finish crowdsourcing tasks, as many crowdsourcing tasks rely on ensemble classification. The other purpose is publishing the learned classifier ensembles as crowdsourcing result. This happens in reciprocal crowdsourcing tasks, where each participant contributes limited data (due to privacy consideration) to the centralized server, to collectively train a ensemble classification model, and then receive the learned classifier ensemble from the server, who has much more computation and storage resources than each single mobile devices for training the ensemble model. Consequently, each participant can then do classification tasks on their own mobile devices without sacrificing data privacy. In concise, during mobile crowdsourcing, the server may need to publish pruned classifier ensembles to plenty of participants that equipped with mobile devices.

We argue that participants need customized levels of accuracy-cost trade-off in ensemble pruning, as some participants may pursue better classification accuracy, while some other may pay more attention to energy saving or time saving during classification. One naive approach to tackle such personalized ensemble pruning challenge is running existing ensemble pruning approaches multiple times with different trade-off parameters, but this would be too costly especially when the number of participants is relatively large.

Motivated by resilience of bi-objective optimization, which simultaneously optimizes for two objectives, in this paper, we embed personalized ensemble pruning problem into bi-objective Pareto optimization problem by taking the classification accuracy and its cost as the objectives. We also propose an improvement in Pareto optimization by mixing up the accuracy and the cost objectives, such that the searching space of ensembles is concentrated when each participant's trade-off level is known by the server. More specifically, this work makes the following contributions:

- We formulate the personalized ensemble pruning problem, which captures the need of personalized accuracy-cost trade-off levels of crowdsourcing participants.

– We present a personalized ensemble pruning framework based on bi-objective Pareto optimization, we further provide an objective-mixture improvement to Pareto optimization for linear-combination loss function.
– We conduct extensive experiments to validate the efficiency and effectiveness of our framework.

The remainder of this paper is organized as follows. Section 2 provides preliminary knowledge about bi-objective optimization and ensemble pruning. Section 3 gives our personalized ensemble pruning model and framework. Section 4 describes the objective-mixture improvement on Pareto optimization. Section 5 theoretically analyzes our framework. Section 6 presents experimental results. Section 7 reviews related work on ensemble pruning. Lastly section 8 concludes the whole paper.

2 Preliminaries

In this section, we briefly introduce bi-objective Pareto optimization and the formal definition of ensemble pruning.

2.1 Bi-objective Pareto Optimization

Multi-objective optimization explicitly optimizes for two or more objectives and outputs multiple solutions, in which each solution has at least one advantage objective when compared to any other solutions. Conventionally the *Pareto optimal* rule is used to determine the advantage or domination relation between solutions, and a solution is Pareto optimal if it is not dominated by any other solutions. Specifically, the Pareto domination relation in bi-objective optimization is formulated as follows:

Definition 1 (Pareto Dominance). *Let $f = (f_1, f_2) : \Omega \to \mathbb{R}^2$ be the objective vector. For two solution $S_1, S_2 \in \Omega$, S_1 dominates S_2 iff $f_1(S_1) \leq f_1(S_2)$ and $f_2(S_1) \leq f_2(S_2)$, and $f_1(S_1) < f_1(S_2)$ or $f_2(S_1) < f_2(S_2)$ holds.*

Unlike one-objective optimization that usually only have one optimal solution, bi-objective optimization returns a set of Pareto optimal solutions. This resilience property of bi-objective optimization lately motivates our framework to use bi-objective optimization to simultaneously approximate optimal classifier ensembles with diverse trade-off levels.

2.2 Ensemble Pruning

Ensemble pruning selects a subset of classifiers S from a pool of base classifiers T, such that the accuracy of aggregated classification results only slightly degrades or even improves, meanwhile the number of classifiers (or computational costs) in S is significantly reduced compared to the whole classifier pool T.

To be precise, let the number of base classifiers in $T = \{t_0, ..., t_m\}$ be m, and selected classifiers $S \in \{0, 1\}^m$ be a boolean vector, in which each bit indicates

whether each t_i is selected. Ensemble pruning seeks for a trade-off between the classification error rate $E(S)$ on a validation data set and the cost to evaluate classifiers in S as follows:

$$S_{opt} = \arg\min_{S\in\{0,1\}^m} E(S) + \alpha \cdot |S|, \tag{1}$$

Where $\alpha \in [0.0, +\infty)$ is the trade-off level, $|S| = \sum_{k=1..m} s_i$ is number of pruned classifiers, and $E(S) + \alpha \cdot |S|$ is the combined loss function. Without loss of generalization, we treat the evaluation cost of each classifier as 1 unit here, but one can also use the actual time or(and) memory consumption of evaluating each classifier on mobile devices as the computational costs.

3 Model and Framework

3.1 Personalized Ensemble Pruning Model

Recall that during crowdsourcing, the server usually need to publish pruned classifier ensembles to crowdsourcing participants. In practical, some crowdsourcing participants may pursue best achievable classification accuracy for better quality of services, while some others may pay more attention to energy saving or time saving during classification. This implies that ensemble pruning with personalized accuracy-cost trade-off levels is needed in mobile crowdsourcing scenarios.

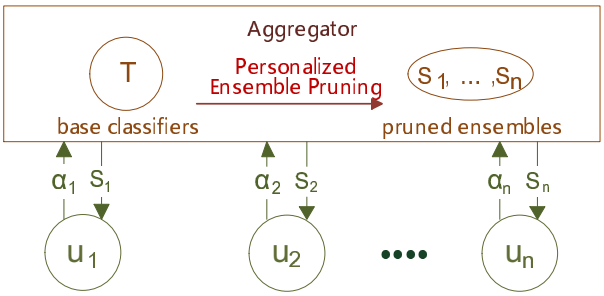

Fig. 1. Publishing personally pruned classifier ensembles in crowdsourcing.

Formally, supposing that there are n crowdsourcing participants, each has a choice of the accuracy-cost trade-off level α_j for ensemble pruning, then the optimal pruned ensemble $S_{j,opt}$ for participant u_j can be formulated as follows:

$$S_{j,opt} = \arg\min_{S\in\{0,1\}^m} E(S) + \alpha_j \cdot |S|. \tag{2}$$

Consequently, the aggregator should publish personalized optimal classifier ensemble $S_{j,opt}$ to the participant u_j, as showed in Fig.1. But practically, as finding one optimal classifier ensemble $S_{j,opt}$ in exponential-sized search space is already computationally hard, a approximation S_j of $S_{j,opt}$ is preferable instead.

The main challenge for personalized ensemble pruning is how to efficiently approximate all $S_{j,opt}$ for each participant when the number of participants is relatively large.

3.2 Personalized Ensemble Pruning Framework

We now provide our basic personalized ensemble pruning framework based on bi-objective optimization. Since finding the optimal pruned ensemble for each participant with trade-off level α_j would be too costly, especially when the number of participants n is large, here we resort to bi-objective Pareto optimization, which takes the empirical classification error $E(S)$ on validation data set as objective $f_1(S)$ and computational costs $|S|$ as objective $f_2(S)$ in Definition 1, and returns a Pareto optimal ensemble set $P = \{p_1, p_2, ..., p_l\}$ as solutions, where l is size of the Pareto optimal solution set.

Algorithm 1 Personalized Ensemble Pruning Framework

Input: A set of trained base classifier $T = \{t_1, t_2, ..., t_m\}$,
participants' trade-off levels $A = (\alpha_1, \alpha_2, ..., \alpha_n)$.
Output: Personalized pruned ensembles $\mathbb{S} = (S_1, S_2, ..., S_n)$ for each participant.

```
Let f(S) = (E(S), |S|) be the bi-objective.
P = bi-objective-solver(f(S))
for j = 1 to n do
    S_j = arg min_{S∈P} E(S) + α_j · |S|
end for
return S = (S_1, S_2, ..., S_n)
```

As detailed in Algorithm 1, in our framework, after a tractable Pareto ensemble set P is returned by bi-objective solver, the approximate personalized optimal pruned ensemble S_j is selected from P. As long as the Pareto optimal set P is sufficiently optimized and complementary, the personalized pruned ensemble S_j for participant u_j will be a good approximation to $S_{j,opt}$ in Formula 2.

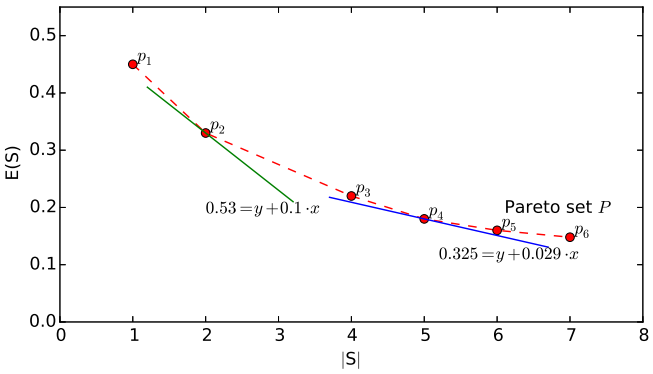

Fig. 2. An example of selecting personalized pruned ensembles from the Pareto set P.

In Fig. 2, we demonstrate how the framework selects personalized pruned ensembles for two participants with trade-off levels 0.1 and 0.029 from Pareto ensemble set P. In this case, the tractable minimum combined loss in Formula 2 for participant with trade-off level 0.1 is 0.534, for participant with trade-off level 0.029 is 0.325.

Naively, selecting optimal ensemble for each participants from Pareto set in line $3-5$ or in the Fig. 2 has computational complexity of $n \cdot l$, which is inefficient when number of participants is relatively large, where l is size of P and is related to number of iterations in the bi-objective solver, hence $l = O(poly(m))$. Fortunately, by sorting the trade-off levels in A and the Pareto set P as showed in Algorithm 2, the complexity can be optimized.

Algorithm 2 Personalized Ensemble Selection

Input: The Pareto ensemble set $P = \{P_1, P_2, ..., P_l\}$,
participants' trade-off levels $A = (\alpha_1, \alpha_2, ..., \alpha_n)$.
Output: Optimal pruned ensembles $\mathbb{S} = \{S_1, S_2, ..., S_n\}$ in P for each participant.

```
sort P by |P_i| in descending order
sort A in ascending order
i = 1
for j = 1 to n do
    while i < l and E(P_{i+1}) + α_j · |P_{i+1}| ≤ E(P_i) + α_j · |P_i| do
        i = i + 1
    end while
    S_j = P_i
end for
return 𝕊 = {S_1, S_2, ..., S_n}
```

Note that any algorithm solves the bi-objective Pareto optimization problem can be deployed in our framework. Recently, Qian *et al.*[13] proposed an evolutionary algorithm for non-personalized ensemble pruning with guaranteed theoretical advantages over former approaches, this outstanding work on non-personalized ensemble pruning can naturally be embedded into our framework for personalized ensemble pruning.

4 Objective-Mixture Optimization

Recall that in our basic framework, the bi-objective $f(S) = (E(S), |S|)$ solver returns a set of Pareto optimal ensembles $P = \{p_1, p_2, ..., p_l\}$, among which for any $p_i, p_j \in P$, p_i doesn't dominate p_j. However, when taking the concrete value of participants' trade-off levels $A = \{\alpha_1, \alpha_2, ..., \alpha_j\}$ into consideration, the non-dominate relation between p_i and p_j may not hold for any level $\alpha \in A$, formally we have:

$$(E(p_i), |p_i|) \not\succeq (E(p_j), |p_j|) \;\nRightarrow\; (E(p_i) + \alpha \cdot |p_i|) \not\succeq (E(p_j) + \alpha \cdot |p_j|). \qquad (3)$$

This implies that the Pareto optimal ensemble set $P = \{p_1, p_2, ..., p_l\}$ still has redundances when concrete trade-off levels $A = \{\alpha_1, \alpha_2, ..., \alpha_j\}$ or their ranges are already known. Consequently, some redundant $p_i \in P$ will not be selected as pruned ensemble by any participants, but will mislead optimization in the bi-objective solving process. In order to mitigate such redundances, we tweak the bi-objective $f(S)$ in line 1 of Algorithm 1 as follows:

$$f(S) = (E(S) + \alpha_{min} \cdot |S|, E(S) + \alpha_{max} \cdot |S|), \tag{4}$$

where $\alpha_{min} = min\{\alpha_1, \alpha_2, ..., \alpha_j\}$, $\alpha_{max} = max\{\alpha_1, \alpha_2, ..., \alpha_j\}$.

As a result, with this optimized objective-mixture framework, for any $p_i, p_j \in P$, there exists an $\alpha \in [\alpha_{min}, \alpha_{max}]$ that the following formula holds:

$$\begin{aligned}&(E(p_i) + \alpha_{min} \cdot |p_i|, E(p_i) + \alpha_{max} \cdot |p_i|) \not\succeq (E(p_j) + \alpha_{min} \cdot |p_j|, E(p_j) + \alpha_{max} \cdot |p_j|) \\ &\Rightarrow (E(p_i) + \alpha \cdot |p_i|) \not\succeq (E(p_j) + \alpha \cdot |p_j|).\end{aligned} \tag{5}$$

Roughly speaking, any $p_i \in P$ in this optimized framework with objective-mixture has the potential to be finally selected as pruned ensemble by a participant.

Practically, when α_{min} and α_{max} is too close, which may cause bi-objective optimization converges local optimally, some additive or multiplicative relaxations could be putted on α_{min} and α_{max}, which means the relaxed objectives is $f(S) = (E(S) + c_{min} \cdot \alpha_{min} \cdot |S|, E(S) + c_{max} \cdot \alpha_{max} \cdot |S|)$, where $c_{min} \leq 1.0, c_{max} \geq 1.0$ are the multiplicative relaxation factors.

5 Performance Analyses

In this section, we analyse the theoretical computational complexity and ensemble pruning performance of our personalized ensemble pruning framework.

Computational complexity of our framework highly depends on the deployed bi-objective solver at line 2 in Algorithm 1. Fortunately, the evolutionary bi-objective solver in [13] shows that expected $O(m^2 \log m)$ iterations is sufficient for approximating optimal Pareto set, where m is the number of base classifiers.

The personalized pruned ensemble selection in line $3 - 5$ naively has computational complexity of $n \cdot l$, but in Algorithm 2, its complexity is reduced to $O(n \log n + l \log l)$, where l is size of P and $l = O(poly(m))$.

To put it concisely, the overall expected computation complexity of our framework is $O(poly(m) + n \log n)$, while the naive ensemble pruning approach for personalized ensemble pruning tends to cost $\Theta(poly(m) \cdot n)$. Hence our framework is efficient in terms of the number of participants n.

The gap between selected pruned ensemble for each participant in our framework and the optimal ensemble in Formula 2 also highly depends on the deployed algorithm of bi-objective solver. The evolutionary bi-objective solver in [13] has been proved to outperform former ensemble pruning approaches, thus the achievable performance of our framework is guaranteed by [13].

6 Experiments

In this section, we evaluate the computational overheads and pruning performance of our personalized ensemble pruning framework, with comparison to the state-of-art ensemble pruning approach PEP in [13], which is also used by our framework as the bi-objective solver at line 2 in Algorithm 1 (the differences between PEP and our framework is declared in the next section).

We mainly compare the running time and combined loss function (in Formula 2) performance of personalized ensemble pruning in our basic framework (BF) and the objective-mixture framework (OMF) with the approach PEPs that runs the PEP multiple times to get personalized pruned ensembles.

In order to simulate ensemble pruning for crowdsourcing participants that equipped with mobile devices, the human activity recognition data set based on mobile phone accelerometers readings from [9] is used for our experiments. The data set has 4944 examples with 43 attributes extracted from time-serial data of mobile phone accelerometers, includes six classes of human activities: walking, jogging, upstairs, downstairs, sitting and standing.

In our experiments, the dataset is randomly split into tree parts: 60% as training set, 20% as validation set and 20% as test set. The pool of $m = 20$ base CART tree [2] classifiers are generated by Bagging [1] on the training set. In these experiments, we simulated with number of participants n ranges from 1 to 1280, with randomized personalized trade-off level $\alpha_j \in A$ ranges from 0.01 to 0.2, and the empirical multiplicative objective relaxation factors in OMF are empirical values $c_{min} = 0.3, c_{max} = 1.7$. The number of iterations of the bi-objective solver is set as 80. We conduct experiments on the scikit-learn [12] platform.

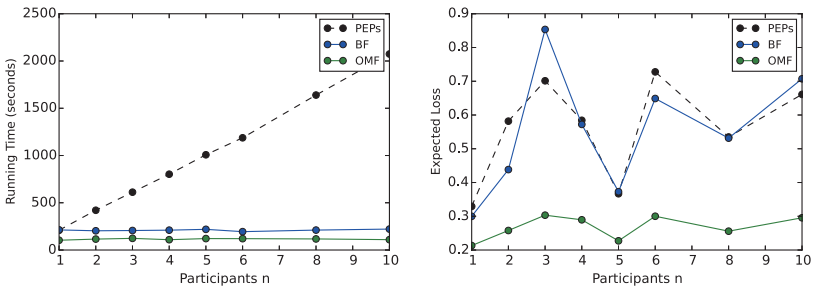

(a) running time (b) expected loss

Fig. 3. The running time and expected loss (the smaller the better) of personalized ensemble pruning for $\{1, ..., 10\}$ participants.

The running time results with a relatively small number of participants are showed in Fig. 3 (a). The PEPs approach for personalized ensemble pruning costs linear to n computational time, while the time consumed by our BF and OMF approach only slightly increase with the number of participants n increases.

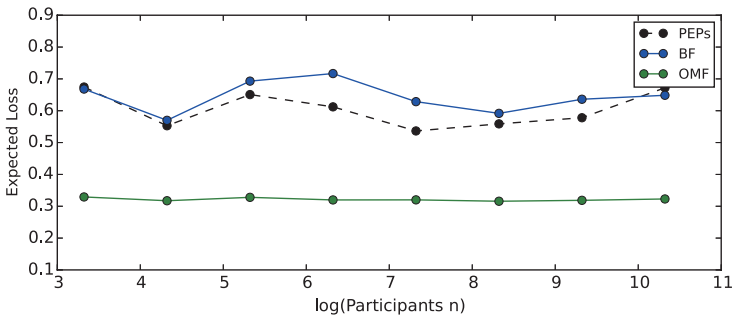

Fig. 4. The expected loss (the smaller the better) of personalized ensemble pruning for $\{10, 20, 40, ..., 640, 1280\}$ participants.

Besides, our framework with objective-mixture (OMF) averagely reduced 40% running time than the basic framework (BF), mainly due to filtering out unnecessary ensemble candidates during each iteration. In concise, our framework for personalized ensemble pruning is efficient when the number of participants n increases.

The expected loss function (in Formula2) results of pruned ensembles for small number of participants are showed in Fig. 3 (b), it demonstrates that there is no significant difference of loss function between the BF and the PEPs approach, while our framework with objective-mixture (OMF) reduced losses by 50% when compared to the BF and PEPs approach. Combining experimental results in Fig. 3 (a) as conclusion, our OMF approach outperforms the PEPs approach and our BF approach in both efficiency and effectiveness.

We also evaluate the performance of our approaches with relatively large number of participants, and the results showed in Fig. 4. The results demonstrate that the losses of pruned ensemble from objective-mixture OMF approach averagely outperforms the basic framework and the PEPs approaches by 50%.

7 Related Work

Ensemble learning is widely used for data classification, in the form of Bagging[1], Boosting[6] or heterogeneous classifiers (e.g. combination of SVM, k-NN and Naive Bayes classifiers). Mainly to mitigate the inefficiency issue in ensemble learning, ensemble pruning has been extensively studied. As ensemble pruning itself is computational hard, genetic algorithms[17] and semi-definite programming[16] has been proposed.

Demir *et al.*[4] and Ulaş *et al.*[15] also consider costs of computations of ensemble or discriminants between classifiers to heuristically compose a pruned ensemble. Some other ensemble pruning approaches such as Margineantu *et al.* [10] and Martínez-Muñoz *et al.* [11] are based on ordered aggregation, which rearrange the order of classifiers in base classifier pool. These approaches pruning ensemble with only one identical trade-off level and are hard modifying for personalized ensemble pruning.

Recently, Qian *et al.* [13] proposed an ensemble pruning approach using evolutionary bi-objective optimization, and proves its theoretical advantages over former heuristic or ordered based approaches. However, it doesn't considers personalized ensemble pruning for crowds nor the objective-mixture optimization when the personalized trade-off levels are known by the centralized aggregator. The concurrent work [5] considers customized accuracy-cost trade-off on a group of Pareto optimal sets, whereas this work focus improving efficiency and effectiveness of personalized trade-off on one Pareto optimal set.

Another topic related to multiple ensembles selection is *dynamic ensemble selection* [8], which dynamically select a classifier ensemble for one class of input samples, instead of for personalized accuracy-cost trade-off levels.

8 Conclusion

In this paper, we formulate the problem of personalized classifier ensemble pruning in mobile crowdsourcing, and proposed a framework for the problem based on bi-objective Pareto optimization with objective mixture. In our framework, each participant (the ensemble user) could choose a customized level of trade-off between the ensemble's classification accuracy and its computational costs, so that the need of energy saving on mobile devices and the need of classification accuracy are balanced for each participant. Both theoretical and experimental results demonstrate the efficiency and effectiveness of our framework. Specifically, the experimental results show that our personalized ensemble pruning framework with objective mixture averagely reduced 40% running time and 50% combined losses simultaneously under same number of iterations.

References

1. Breiman, L.: Bagging predictors. Machine learning 24(2), 123–140 (1996)
2. Breiman, L., Friedman, J., Stone, C.J., Olshen, R.A.: Classification and regression trees. CRC press (1984)
3. Catal, C., Tufekci, S., Pirmit, E., Kocabag, G.: On the use of ensemble of classifiers for accelerometer-based activity recognition. Applied Soft Computing 37, 1018–1022 (2015)
4. Demir, C., Alpaydin, E.: Cost-conscious classifier ensembles. Pattern Recognition Letters 26(14), 2206–2214 (2005)
5. Fan, L., Xue, M., Chen, S., Xu, L., Zhu, H.: Poster: Accuracy vs. time cost: Detecting android malware through pareto ensemble pruning. In: Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security. pp. 1748–1750. ACM (2016)
6. Freund, Y., Schapire, R.E.: A decision-theoretic generalization of on-line learning and an application to boosting. Journal of computer and system sciences 55(1), 119–139 (1997)
7. Gibson, R.M., Amira, A., Ramzan, N., Casaseca-de-la Higuera, P., Pervez, Z.: Multiple comparator classifier framework for accelerometer-based fall detection and diagnostic. Applied Soft Computing 39, 94–103 (2016)

8. Ko, A.H., Sabourin, R., Britto Jr, A.S.: From dynamic classifier selection to dynamic ensemble selection. Pattern Recognition 41(5), 1718–1731 (2008)
9. Kwapisz, J.R., Weiss, G.M., Moore, S.A.: Activity recognition using cell phone accelerometers. ACM SigKDD Explorations Newsletter 12(2), 74–82 (2011)
10. Margineantu, D.D., Dietterich, T.G.: Pruning adaptive boosting. In: ICML. vol. 97, pp. 211–218. Citeseer (1997)
11. Martínez-Muñoz, G., Suárez, A.: Pruning in ordered bagging ensembles. In: ICML. ACM (2006)
12. Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al.: Scikit-learn: Machine learning in python. The Journal of Machine Learning Research 12, 2825–2830 (2011)
13. Qian, C., Yu, Y., Zhou, Z.H.: Pareto ensemble pruning. In: AAAI. pp. 2935–2941 (2015)
14. Sun, Z., Tang, S., Huang, H., Huang, L., Zhu, Z., Guo, H., Sun, Y.e.: iprotect: Detecting physical assault using smartphone. In: Wireless Algorithms, Systems, and Applications, pp. 477–486. Springer (2015)
15. Ulaş, A., Semerci, M., Yıldız, O.T., Alpaydın, E.: Incremental construction of classifier and discriminant ensembles. Information Sciences 179(9), 1298–1318 (2009)
16. Zhang, Y., Burer, S., Street, W.N.: Ensemble pruning via semi-definite programming. The Journal of Machine Learning Research 7, 1315–1338 (2006)
17. Zhou, Z.H., Wu, J., Tang, W.: Ensembling neural networks: many could be better than all. Artificial intelligence 137(1), 239–263 (2002)
18. Zhu, Y., Zhong, E., Lu, Z., Yang, Q.: Feature engineering for semantic place prediction. Pervasive and mobile computing 9(6), 772–783 (2013)